



P1

P2

P3

P4

P5

P6

P7

P8

T1

T2

T3

T4

T5

T6

M1

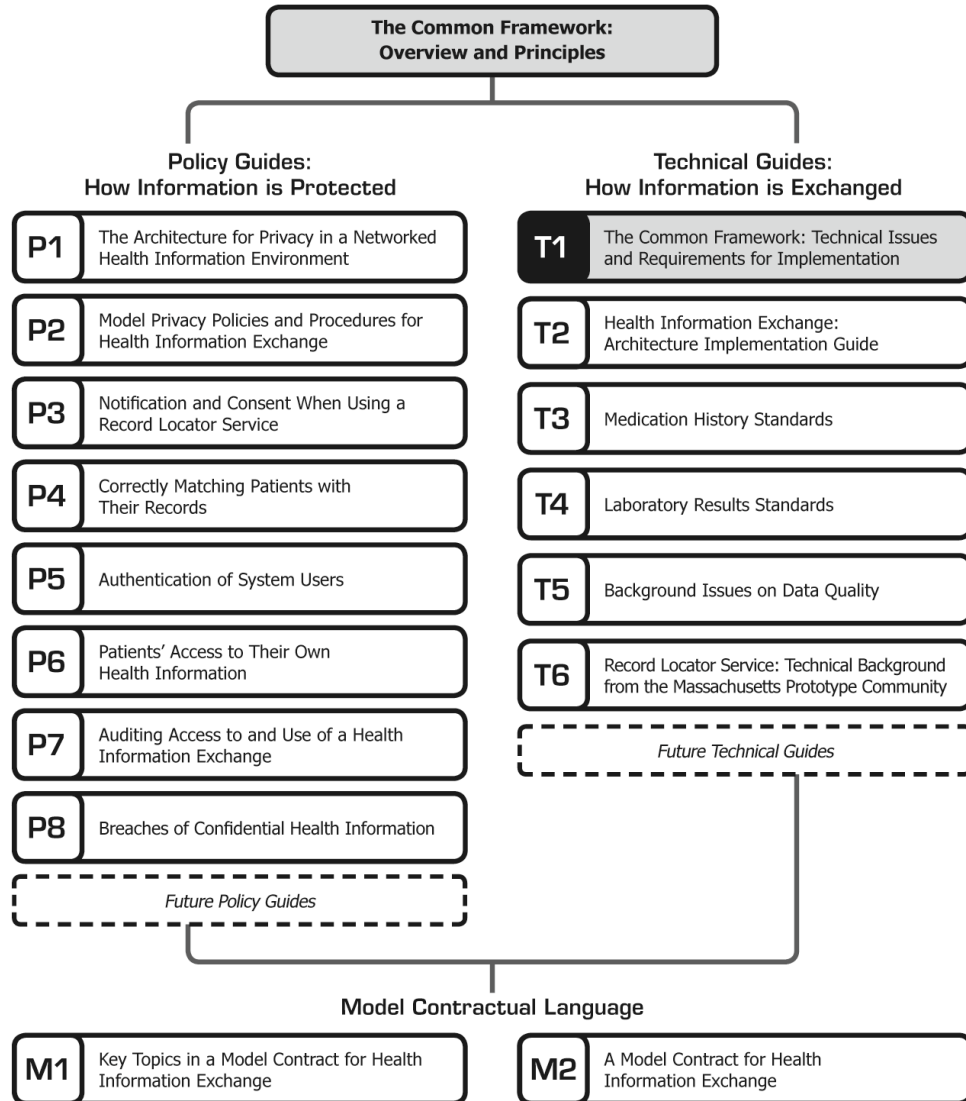
M2

The Common Framework:

Technical Issues and
Requirements for Implementation

The Common Framework: Technical Issues and Requirements for Implementation

The document you are reading is part of *The **Connecting for Health** Common Framework*, which is available in full and in its most current version at: <http://www.connectingforhealth.org/>. The Common Framework will be revised and expanded over time. As of April 2006, the Common Framework included the following published components:



The Common Framework: Technical Issues and Requirements for Implementation*

What is this document?

This document is an overview of the technical philosophy and decisions behind **Connecting for Health's** Common Framework, and to the policy issues related to those technical requirements. It describes, in broad outline, a vision for a nationwide health information network (NHIN) that preserves patient privacy while allowing health information to be accessed by authorized persons; that leverages the existing investments in health care IT that have been made by existing institutions; and which preserves a high degree of both authority and autonomy for the institutions that currently provide care. It is not a technical guide; rather it is a general overview of the issues, accompanied by pointers to the relevant policy and implementation documents.

In particular, this document provides an outline of the technical issues and choices implicit in creating collaborative networks of health care providers. We call these collaborative networks sub-network organizations (SNOs) because they are components of the larger nationwide network. A SNO will adopt and contractually enforce local standards and policies among its members, and adopt standards and policies that will allow it to inter-operate with other SNOs. The NHIN is simply a network of these SNOs; there is no centrally managed database or set of services within the NHIN separate from those provided by the SNOs themselves.

This document is part of *The Connecting for Health Common Framework: Resources for Implementing Private and Secure Health Information Exchange*, and is accompanied by several other companion technical and policy guides for health information exchange. These documents are intended for anyone in the health care or technology industries interested in health information exchange, but, taken as a sequence, represent various levels of increasing technical detail and complexity.

This document, "The Common Framework: Technical Issues and Requirements for Implementation," provides the most basic level of technical overview. The "Health Information Exchange: Architecture Implementation Guide" covers a deeper level of detail, and the XML files and source code used in the technical prototype test, available through <http://www.connectingforhealth.org>, provide a further level of technical detail. These latter files have been provided only to demonstrate the code that a technical developer created to implement specifications.

Implementation of the technologies described herein assumes that the health care entities engaging in health information exchange have already made significant strides using health information technology (HIT) locally. This requirement does not refer to the complexity or sophistication of the technology, but rather the basic capabilities presumed by the following:

- The participating entities can receive and use digital clinical data, e.g., information brought in from remote sources such as laboratories.

* **Connecting for Health** thanks Clay Shirky, Chair, Technical Subcommittee, and Adjunct Professor, New York University Graduate Interactive Telecommunications Program, for drafting this paper.

©2006, Markle Foundation

This work was originally published as part of *The Connecting for Health: Resources for Implementing Private and Secure Health Information Exchange* and is made available subject to the terms of a license (License) which may be viewed in its entirety at: www.connectingforhealth.org/license.html. You may make copies of this work; however, by copying or exercising any other rights to the work, you accept and agree to be bound by the terms of the License. All copies of this work must reproduce this copyright information and notice.

- The participating entities are already in compliance with HIPPA and state requirements governing data privacy and security, and are capable of implementing the requirements for providing only authorized access to identified individuals within their enterprise or organization.
- The participating entities are comfortable disambiguating patient identities using probabilistic matching algorithms within their enterprise or organization. (Comfortable solving the "John Smith" problem in large databases, in other words.)
- The participating entities have or are willing to acquire the hardware, software, and technical expertise necessary to support secure information exchange over the Internet, using Web Services standards and Secure Socket Layer (SSL) certificates.
- The participating entities are willing to collaborate on the design and implementation of standards and policies for health information exchange among themselves, establish or identify an existing entity to host the Record Locator Service and put in place a governance model consistent with the policy principles.

Part I: Understanding the NHIN

The **Connecting for Health** Common Framework describes and defines relationships between three kinds of organizations: individual health care entities, groups of those entities joined together to form a sub-network organization (SNO),¹ and the nationwide health information network (NHIN) as a whole.

Entity: An entity refers to functionally independent participants within the health care system, from single doctor practices up to hospital chains and national organizations, whether public or private.² An entity is any organization, institution, or practice; the only parts of the health care system that are not entities are the patients themselves.

SNO: A SNO is any group of entities (regionally or non-regionally defined) that agree to communicate clinical data with one another using a single Record Locator Service (RLS), using shared policies and technological standards, and operating together under a single SNO-wide set of policies and contractual agreements. A SNO has two sets of interfaces, one internal, which binds its member entities together, and one external, which is where traffic to and from other SNOs and outside entities come from.

NHIN: The NHIN is the sum of all SNOs. It is a network of networks whose participants agree to the Common Framework. The NHIN is not a separately funded entity; it is a framework of cooperation and compliance.³ If the individual SNOs externally facing interfaces work, the NHIN will work. There are no required "top level" services in the NHIN; at the national level, adherence to standards and policies, however defined and affected, are the key elements. All the actual infrastructure of the network is either hosted within the SNOs, or uses the existing internet.

The following diagrams illustrate the relations between entities, the SNO, and the NHIN.

¹ Another common acronym for this type of organization is RHIO, which stands for Regional Health Care Information Organization. Though a SNO is conceptually similar to a RHIO, we use SNO because there are a number of national or supra-regional institutions such as the VHA, the CDC, health plans, pharmacy chains, and State regulatory agencies that are not defined by regional boundaries and that need to connect with entities in more than one geographic region. Unlike the regional focus of RHIO, a SNO is any group of data-sharing entities that agree to be bound contractually by technical and policy standards, regardless of actual geographic proximity.

² Though we chose the word entity for its obvious parallel to the definition of 'covered entities' under HIPAA, there may be entities in a SNO that are not covered under HIPAA, such as data centers. These entities nevertheless need to comply with the Common Framework, or the entities that employ them need to agree to take on the responsibility of ensuring that compliance.

³ This is sometimes called the 'thin NHIN' model; it assumes a high degree of autonomy and control remains with today's health care information providers, and that no significant new national technical organization is required to 'operate' the NHIN as a whole. Instead, the existence of policies, standards, and connectivity allows the secure sharing of data with authorized persons nationwide without staffing new sites of central management or control.




Diagram 1: Inter-communicating entities within a SNO.

A SNO is a collection of clinics, hospitals, labs, and other entities. They can communicate directly with one another, and are all governed and contractually bound by SNO-wide policies. With the exception of Common Framework technical standards and policy requirements, SNOs are free to design the organizational form that best fits their members. For instance, some SNOs will want to offer a number of SNO-hosted technical services such as data validation, while others will want those functions handled by the member entities.




Diagram 2: Two SNOs communicating securely over the Internet.

Every SNO must maintain an interface to other SNOs. This interface is called the Inter-SNO Bridge (ISB.) As the name implies, the ISB is the point of contact between SNOs. While data traffic within a SNO can pass directly between member entities, traffic between entities in different SNOs must pass through an ISB. The ISB exists both to simplify nationwide traffic (so that every entity does not have to know the address of every other entity), and to improve security, by providing a single point of observation for data traffic going to and from the SNO.




Diagram 3: The NHIN is the sum of intercommunicating SNOs.

The NHIN is the sum of all such intercommunicating SNOs, plus those entities required to design, promote and enforce those standards and policies that need to be implemented at a national level. As the nation's health care institutions join together in SNOs, in order to increase the quality of local care while reducing costs associated with incomplete or duplicated data, the NHIN grows as a side effect.

Creation of a SNO

The SNOs are the principal site of implementation of the Common Framework. The size and scope of a SNO is variable; the commonest model for a SNO is likely to be formal incorporation of a group of entities that have already frequent dealings with one another. The principal function of a SNO is to provide a way to share information securely and among authorized users while protecting patient data against accidental or inappropriate disclosure or misuse. The SNO will contractually mandate and ensure compliance with Common Framework standards and policies, and will adopt or enact any additional standards or policies it deems necessary, with the consent of its members.

Levels of SNO Guidelines

There are four conceptual levels at which a given policy or technical standard can be set: national, SNO-wide, per member institution, or no standard.

	Example
National	Encryption standards for sending data securely
Per SNO	Acceptable fields for patient ID
Per Entity	Management of user ID and authentication
No Standard	Clinical applications and interfaces

In keeping with a desire to make the Common Framework as broadly adoptable as possible, we have tried to specify the minimum necessary, assuming that above some floor, different entities will have the flexibility to do more; and we have tried to specify policies and standards at the appropriate level of adoption and enforcement.

For example, we must have a national standard for encryption of data as it passes between SNOs, both to assure that the data is secure and to provide interoperability. Coordination of acceptable identifiers for patients in the absence of a national identifier, however, cannot be handled nationally, because different states regulate e.g., the use of Social Security Numbers (SSNs) differently. This level of standardization must therefore be done by the SNO. Procedures for handling user identity and authentication must be in place, but we cannot require that they be handled at the level of the SNO, because in many cases, members of a SNO will have existing procedures which they will not be able to replace quickly or easily. In these cases, the policy and technical documents set minimum levels of compliance, but do not require SNO-wide standardization.

Lastly, and most importantly, there are those aspects of the health care system that the Common Framework does not propose to standardize. For example, we do not propose standardizing the 'look and feel' or behavior of clinical applications, because such standardization would stifle innovation and is not necessary for different applications to share data (so long as the data itself is in a standard format). The design and implementation of clinical applications is best served by allowing those applications to be tailored to local needs and requirements -- a well-funded research hospital will have different application requirements than a safety-net clinic with only web browser access to remote data.

Our view has been that where it is not necessary to standardize, it is necessary not to standardize. Though it is tempting to believe that one can upgrade the entire health care system at once, experience has shown that every required standard raises the cost and difficulty of complying, while lowering the overall rate of adoption. In a system as large and diverse as this one, the essential tasks of improving data quality and creating applications better able to take advantage of that data must be approached incrementally. We believe that the best way to catalyze those improvements are to decisively separate the functions of the network for discovering and transporting that data from the applications that request and consume it.

Part II: Applications and Interactions

Modeling Interactions

The relationship between health care entities in the NHIN can be understood by imagining five sample interactions:

1. Transfer of patient's clinical data between entities in a single SNO
2. Between entities in two different SNOs
3. From an entity in a SNO to an unaffiliated entity
4. From an entity in a SNO to a public health entity or other aggregator
5. Subscription Models for the above 4 Interactions

1) Transfer of Clinical Data Between Two Entities in a Single SNO

Step 1: Asking for Record Locations

A patient, Elizabeth Smith, presents at a clinic complaining of shortness of breath. She has never been seen there before, and after she provides basic demographic data, the clinic queries the local RLS for her records.⁴ Assuming the clinic is a member of the SNO, and has presented the proper credentials for authentication and authorization, the RLS will compare a set of identifiers, for example Ms. Smith's name, DoB, gender, Zip, and SSN, with those records whose locations are listed in the RLS database. Those record locations that have a sufficiently high probability of matching the patient data (as determined by a matching algorithm specifically designed for that purpose) will be returned from the RLS.

The RLS answers all queries from authorized sources within a SNO.

Step 2: Aggregating the Identified Records

Once the RLS has matched Ms. Smith's identifying details against record locations it contains in its database and returned information necessary to retrieve those records, its work is done. The next step is to use those locations to aggregate the actual records themselves, by querying the record locations for the actual patient data. Because the site of the aggregation has significant ramifications for the contractual relations within the SNO the site of aggregation of records can vary SNO by SNO. There are several options for aggregating the actual records whose locations have been returned by the RLS. The **Connecting for Health** prototype tried three; others may be possible.

Client Aggregation: One method is to have the original requesting client do the aggregating. In this model, the client receives one or more record locations. (Zero locations returned is a failed search.) The client then decides which of these records it would like to attempt to retrieve, e.g., only records from a particular institution, or only records from labs, or all records.⁵ Client-side aggregation was tested in Massachusetts. The advantages of client aggregation are refined control over record requests, and possibly tighter integration with other local electronic data systems. The disadvantage is higher technical requirements at the participating entities in the SNO.

Central Aggregation Service: A second method of aggregation is to create a central aggregation service, a server that sits between the entities and the RLS itself. The server takes incoming queries and passes them on to the RLS, and then takes the record locations returned and queries each listed location, handing off only the final, aggregated record to the requesting client. Centralized aggregation was tested in Indiana. The advantages of such a service are that it creates economies of scale for the SNO. The disadvantages are less control of the record by the original requestor, and higher security risk, since the aggregation service will hold, even if only for a moment, considerable clinical data about the patient.⁶

Aggregation Proxy: A third possibility is to run a proxy service that can receive record locations and aggregate them, but is not a required service, allowing some clients to use the remote server for aggregation, and others that want to receive the record locations and do the aggregation locally to do so. The proxy aggregation model was tested in California. The advantages of such a system are that it allows aggregation or

⁴ The application that does the querying could be as simple as a secure web browser, or as complex as an integrated medical record system.

⁵ The decision about which of the records to request can be done either with human intervention or by an automated process.

⁶ It is critical, in fact, that any such aggregation service not cache the clinical data it is handling, so that it doesn't become a significantly attractive hacking target.

records to happen either centrally or at the requesting site. The disadvantage is that it potentially more complex, in terms of interaction, than the other two models.

In all three aggregation scenarios, some sort of authentication and authorization must be in place between the requestor and the source of the data, whether peer-to-peer (each entity authenticates directly to each other entity) or, more likely, through the operation of a SNO-wide directory service that allows the entities to identify one another. (See the section on Identity, Authentication, and Authorization, below.)

Importantly, under the Common Framework the actual sources of clinical data are *not* required to respond to any given request for that data. Individual entities are the stewards of the records, and of the patient's expectation of confidentiality. As a result, those entities may add constraints on data access.⁷ Examples include added restrictions to a particular set of records at the patient's request, or added restrictions for anyone who does not have admitting privileges at a particular hospital.

Step 3: Displaying or Otherwise Using the Records

The third step involves taking the actual clinical records, wherever aggregated, and making them useful, whether by displaying them directly to the clinician, integrating them into an existing electronic health record (EHR) system, feeding them into a decision support tool, or any of a number of other present or future possible uses of the data.

A key aspect of the current model is that it places no constraints on how the records are made useful, other than to require that the consuming applications abide by policy requirements around privacy, security, auditing, etc. The role of the network is to carry useful data from existing sources to authorized requestors; whether that data is then displayed directly in a browser window or becomes part of a complex database transaction is entirely up to the local user.⁸ The goal of the Common Framework is to advance the conversation between application designers and users, by making data more accessible and better formatted. The Common Framework is not intended to either replace or interfere with those conversations.

2) Transfer of Clinical Data Between Two Entities in Different SNOs

A similar scenario can occur when a clinic in one SNO (A) looks for a patient's records in a second SNO (B) of which the clinic is not a member. The basic three-step transaction is the same, with these differences:

Step 1: Asking for Record Locations

1. In addition to the patient's demographic details, the clinic needs some information on which other SNOs to query, whether the name of an institution, an affiliation with a particular network (e.g., the VHA), or a region where she previously received care. There

⁷ Uniform response to intra-SNO requests for records are not required, nor are they forbidden. An individual SNO can, as a matter of policy, mandate such uniform access if it decides to do so. This requirement would override the ability of individual institutions to differ in data access policies.

⁸ This design pattern is known as the End to End principle.
<http://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf>

is no national index of patients; the reasons for this are discussed in the ISB section below.⁹

2. All traffic leaving SNO A goes out through SNO A's ISB -- the clinic does not make remote requests directly.
3. Once a request is validated by the receiving ISB, that request is forwarded to the RLS in SNO B and handled as are internal requests, except that the response goes back to the ISB of SNO B for return to SNO A.
4. All traffic coming into SNO B from entities that are not members of B comes in through B's ISB. This simplifies contact with the outside world, and provides a single spot for watching remote traffic (which has a lower level of trust than local traffic.)
5. The trust model between SNOs specifically assumes that each SNO's ISB has a valid SSL certificate, and each SNO will accept the other's certificate.
6. The requesting SNO must provide an identifier of the person authorizing the request. (See the policy document, "Authentication of System Users.") The receiving SNO does not need (and will in most cases be unable to) re-authenticate the original requestor.

Step 2: Aggregating the Identified Records

1. The clinic can, optionally, ask either for a set of pointers to the data, or can ask the ISB to act as an aggregator, and return the aggregated record directly. The ISB must support both types of request.
2. Inter-ISB communications are always asynchronous. ISB A passes along the clinic's request for data to ISB B. B responds that it has received the request. When it is time to return data to A, it starts a transaction with A to deposit that data. Each ISB must therefore be able to both initiate outbound requests to other ISBs and to accept other transactions from ISBs.
3. It is up to SNO A to determine how the material is to be transferred from its ISB back to the initial requesting entity. The ISB can require the original requestor to check back periodically; can maintain an open connection via streaming until the data returns from ISB B; can even email or fax the data if those methods are supported.

Step 3: Displaying or Otherwise Using the Records

As with records received from within the SNO, the current model is that it places no constraints on how the records are made useful, other than to require that the consuming entity abides by policy requirements around privacy, security, auditing, etc

3) Transfer of Data from an Entity in a SNO to an Unaffiliated Entity

A similar scenario can occur when an entity A, which is not part of any SNO, requests a patient's record held in SNO B.¹⁰ This scenario is critical for the network to grow organically, since the early days of any such network will necessarily cover only a minority of potential participants. The basic three-step transaction is the same as the transfer of data between SNOs, with these differences:

Step 1: Asking for Record Locations

⁹ There is nothing to prevent a group of SNOs, for example all SNOs in a given region or state, to develop independently their own inter-SNO index of patients.

¹⁰ It is possible to imagine an entity in a SNO requesting data from an outside entity unaffiliated with a SNO, the reverse of this transaction, but such a transaction would be completely ad hoc, as it involves a data-holding entity ungoverned by the Common Framework.

The trust model between unaffiliated entities and SNOs assumes that any SNO accepting queries from unaffiliated entities will subject such requests to a high standard of scrutiny, and to higher levels of audit, and will in any case not automatically honor such requests without some form of scrutiny.

Step 2: Aggregating the Identified Records

Communication from an ISB to any outside entity is always asynchronous. As a result, any clinic asking for material through an ISB must either have an accessible online receiver for the results, or must have access to a third-party service that offers such a receiving service. The Common Framework is designed to allow for the creation of such third-party services, though in all cases, the sending and receiving parties are responsible for care of the patient's data, and will be liable for any loss occurring through third-party services they hire or subscribe to.

Step 3: Displaying or Otherwise Using the Records

As above, the current model is that it places no constraints on how the records are made useful, other than to require that the consuming entity abides by policy requirements around privacy, security, auditing, etc.

4) From an Entity in a SNO to a Public Health Institution or Other Aggregator

The 2004-2005 work on the Common Framework concentrated on clinical data. However, in addition to handling identified clinical records about individual patients, there are many reasons to handle aggregate and partially anonymized records, including satisfying public health reporting requirements, quality reporting, and fraud detection. This scenario is quite different from any transfer of clinical data, and is handled differently in the Connecting for Health model.

Because the Record Locator Service contains no clinical data, aggregate and anonymized requests are not dispatched directly to the RLS, but instead to the individual institutions, which reply with those requests directly. It is currently up to the individual SNO, in negotiation with the entities who are allowed access to aggregate or anonymized data, to determine whether such requests should go through the ISB or should be handled as direct connections between the entities and the aggregators of the data. This allows the partitioning strategy for protection of data to continue to operate even when handling aggregate data, even when such requests are not governed by HIPAA, as with required public health reporting. Our model for direct aggregation from the source is the Shared Pathology Informatics Network (SPIN),¹¹ and modeling of SPIN-style interactions in the Common Framework is part of the 2006 effort.

5) Subscription Models

Any of the above transactions may be modeled as a subscription to a particular source of data as well, where an authorized user can request that when a piece of remote content is updated, they receive either a notification of the update, or receive the updated data itself. However, this pattern is not yet specified. The Common Framework is based on Web Services, which enormously lowers the required coordination among network participants, both in advance of and during a transaction. As a result of this loose coupling, subscription models of data transfer (e.g., "Notify me when this patient has new lab results.") can be modeled in two ways. The first is 'scheduled pull' -- scheduled requests from the querier to the RLS or data holder, which requests automatically repeat periodically, in intervals between seconds and months

¹¹ <http://www.cancerdiagnosis.nci.nih.gov/spin/>.

depending on the nature of the query. The other is 'triggered push', where the RLS or data holder watches for updates to data, and pushes out any such updates to a list of subscribers or their designated proxies.

The design and implementation of such models is complex and highly dependant on the technical savvy of the member entities of the SNO. A number of variables affect decisions about subscriptions, such as who assumes the costs of maintaining the subscription information (the querier, in the case of scheduled pull, and the holder of the data, in the case of triggered push.) As a result, like aggregation, the design and implementation of subscription models is currently envisioned as a per-SNO design choice, though with the assumption that observation of the various implementations in 2006 will provide a guide to any nationwide standardization.

Broad Policy and Technical Requirements

The Common Framework provides a list of the minimal set of policies and standards that must be adopted by any participating SNO. On the policy and governance side, all incorporated members of a SNO¹² must:

1. Adopt the policies of the Common Framework (See the policy documents contained in *The **Connecting for Health** Common Framework: Resources for Implementing Private and Secure Health Information Exchange.*)
2. Agree to any SNO-wide policies in place

In addition, each SNO has three technical services it must offer:

1. A SNO-wide Record Locator Service, to allow authorized entities within the SNO to look for patient data
2. A matching algorithm, to match patient demographics contained in incoming requests with the records stored in the Record Locator Service.
3. An Inter-SNO Bridge (ISB), to allow authorized outside parties to look for and retrieve patient data

The basic rationale behind these governance and technical requirements are discussed below; the detailed policy recommendations are contained in *The **Connecting for Health** Common Framework: Resources for Implementing Private and Secure Health Information Exchange*; the detailed technical implementation guides are contained in the "Health Information Exchange: Architecture Implementation Guide", contained in *The **Connecting for Health** Common Framework: Resources for Implementing Private and Secure Health Information Exchange*.

A health care entity can belong to more than one SNO; this would of course entail the additional expense of listing patient demographics and record location information in more than one place, and reconciling contractual requirements where they differ between SNOs. There is no conceptual obstacle to multi-SNO membership, however. There is no minimum or maximum size for a SNO; a single institution can be a SNO so long as it adheres to the principles and standards of the Common Framework. In practice, only very large institutions will do this, as having a single institution as a SNO creates little of the efficiencies or cost-savings that multi-entity SNOs can have.

¹² Note that patients may in some cases be considered 'members' of a SNO, if they access their data from interfaces supported by the SNO. These patients are not covered by Common Framework requirements, as they enjoy a different degree of control over their data than incorporated entities do.

Software Requirements: RLS, Matching Algorithm, ISB

One of the key design principles of the Common Framework is that no particular software application is required; in the same way that email software from different organizations all read the same email data standards, the technical infrastructure of a SNO can be built on any suitably secure hardware and software platform,¹³ so long as it produces and consumes common data standards.¹⁴

The three applications a SNO is required to host¹⁵ are the Record Locator Service, a matching algorithm for matching queries for clinical data with patient records, and an Inter-SNO Bridge, for traffic between the SNO and the outside world.

RLS

One of the basic software requirements of a SNO is the operation of a Record Locator Service (RLS.) The institutions with the right to list patient demographics and record locations in the RLS are the members of a SNO, by definition. Thus the RLS is the practical locus of most SNO-wide activity. The details of the RLS are covered in the "Health Information Exchange: Architecture Implementation Guide," and the relevant policies in *The Connecting for Health Common Framework: Resources for Implementing Private and Secure Health Information Exchange*, but the basic functions are described here.

The Common Framework makes the following assumptions about the design of the RLS:

1. There is one RLS per SNO, which holds the universe of records that can be queried using the RLS service within that SNO.¹⁶ There is no meta-RLS, in keeping with the "No requirement for national services" design.
2. The RLS is designed *only* for patient-centered queries. Aggregate queries (e.g., "Find all admissions in the last 24 hours presenting with shortness of breath") must be dispatched to the participating institutions, or run against aggregated databases that are collected and kept separately. The lack of clinical data at the RLS keeps the RLS from being a target of loss or theft of clinical data, and allows interactions to be optimized for a single, simple case.
3. The RLS participates in two types of transactions -- the addition, modification, or deletion of listed patient record locations from the entities that hold data on the patient, and requests for information about a particular patient from entities that want those locations.
4. All transactions to and from the RLS are logged and audited.
5. The RLS must have a valid SSL certificate, and may only communicate with requestors who support encrypted web communications (https).
6. The RLS is designed to take a query from authorized users in the form of demographic

¹³ During the proof-of-concept testing in 2005, interoperating systems used a mix of platforms and tools, including the .Net framework (1.1) running on Windows XP servers and Java application and Web Services servers (Tomcat and Axis) running on Linux.

¹⁴ Note that these requirements don't foreclose optional support of additional data types or interaction patterns. They are simply minimum requirements, so that new entities joining a SNO have an obvious and small set of required standards to support, and so that an entity that wants to belong to more than one SNO is not caught by requirements to support multiple standards.

¹⁵ Because identity, authentication, and authorization services are already required by HIPAA, we treat them separately below

¹⁶ Note that there may be records held by institutions within the SNO that exist but are not listed in the RLS, because of some institutional or patient preference for keeping them unlisted. This is an option for entities in the SNO unless SNO-wide policy overrules such a thing, and is not itself overruled by another regulation such as a state-wide requirement.

- details or, alternatively, a query in the form of a unique provider ID plus a Medical Record Number, which would enable them to use the RLS to find other records for the patient whose MRN they know.¹⁷
7. The RLS *must* support patient data in incoming queries expressed in the HL7 2.4 format described in the "Health Information Exchange: Architecture Implementation Guide."
 8. The RLS *may* support incoming queries expressed in the HL7 3.0 format described in the "Health Information Exchange: Architecture Implementation Guide."
 9. The RLS must support both synchronous queries, where the data is returned in a single round trip, and asynchronous queries, where the data is delivered in a new session, some time after the original query. The querier may request either synchronous or asynchronous queries; the RLS may also default to asynchronous return of results if it is unable to complete a given query in a timely fashion.
 10. The RLS must implement a probabilistic matching algorithm for patient queries so that the chance of incidental disclosure (presenting a false match) is minimized. (See the policy document "Correctly Matching Patients with Their Records.")
 11. In responding to such queries, the RLS will return zero or more matching demographic records, each including a locator (usually an Institution code and a Medical Records Number) to a set of clinical data for that patient. The locator contents are used for subsequent queries for clinical data.
 12. The RLS will return only records which meet or exceed a minimum probability level. (See the policy document "Correctly Matching Patients with Their Records.")
 13. The RLS will not provide a "Break the Glass" procedure in which a physician or other inquirer can request an emergency exception to allow examination of records below the minimum probability level. Besides having a high probability of incidental disclosures and false positive matching, there is no logical additional method that the inquirer can use to positively identify the correct record. If a user has certainty that a record related to a specific patient exists at a particular entity, that user should work directly with that entity to attempt to locate the record.
 14. The RLS will return "as matched" data for any data transformations it performed in matching the data (e.g., noting that it matched a name provided as Elizabeth with a patient whose first name is listed as Liz.)¹⁸
 15. The RLS should not return demographic data in fields not submitted by the querier. The RLS may well have demographic details about a patient that a clinician has not submitted; these details should not be displayed, to avoid incidental disclosure, and the risk of authorized users fishing for data.
 16. The SNO must maintain a logical separation of clinical from demographic (identifying) data. The RLS itself will not hold clinical data or metadata; all of that is controlled by the entities that created the data, or who hold copies because they provide the patient with care.
 17. The design of the RLS assumes that the clinical data itself may be served from cached or other copied versions of the "live" clinical data, and it is acceptable to centralize the *physical* storage of this data, in order to control costs and guarantee service levels. However, wherever it is located, the data itself should remain in the control of the providing institution, which should be deferred to as the final source of truth on issues of data accuracy and cleanliness.
 18. At the time or shortly after records are published to the RLS from entities, the RLS must report obvious errors back to the publishing entity. Such errors include but are not limited to non-numeric characters or incorrect number of digits in numeric data such as SSN, day, month and year designations in Date of Birth; dates that are out of range

¹⁷ This second format assumes that the querier has some method of obtaining such a record number, either because the patient has provided it or because the querier has used it in the past for the same patient.

¹⁸ One caveat is SSN; if a system uses SSN for matches, the RLS should never return the SSN, even if the requestor supplied it, because of the sensitivity of that particular field.

- (e.g., February 31); etc. The requestor is not required by the Common Framework to act on these reports, but the RLS must make them available, and the individual SNO may have a policy requiring particular responses to such errors.
19. At the time of publishing records to the RLS from an entity, the RLS may report possible errors, including but not limited to name and gender fields with a high probability of inconsistency (e.g., Sylvia, M), pairs of records above the matching threshold with different dates of birth; patient records above the matching threshold with different local record numbers, etc. The publishing entity is not required by the Common Framework to act on these reports, but the RLS must make them available, and the individual SNO may have a policy requiring particular responses to such errors.
 20. The RLS must be able to provide an audit log indicating all entities that have published records on behalf of an individual patient and all users that have received record locations in response to requests regarding an individual patient.

Adoption of a Matching Algorithm

The RLS stands between authorized queriers (either entities within a SNO, including possible aggregation services, or the ISB) and a database of patient demographics and record locations. The RLS's job is to take the incoming queries, format the contents of the message, and make a query to a matching algorithm that determines which records in the database are likely matches. Those records, and only those records, are returned by the matching algorithm to the RLS. The policies governing the matching algorithm are covered in "Correctly Matching Patients with Their Records."

There is not any standard matching algorithm that can be adopted nationwide, because the work on matching is highly sensitive to local regulations, as with regions that forbid the use of Social Security Numbers (SSNs) for matching and to the relative "cleanliness" of the underlying data. The more accurate the collection and storage practices are, the more likely that highly accurate matching will occur with fewer fields. (See "Background Issues on Data Quality" in *The Connecting for Health Common Framework: Resources for Implementing Private and Secure Health Information Exchange* for a discussion of this issue.)

Such algorithms are also highly sensitive to local characteristics of the data set being queried. A last name of Smith is a better predictor of uniqueness in Wewoka, OK than Waltham, MA; NYSIIS-transformed¹⁹ names are better matches for Anglo-Saxon names than French names; and so on. The adoption of a matching algorithm that satisfies the conditions below is a nationwide requirement; the nature and tuning of the particular algorithm must be left to the SNO itself.

The Common Framework makes the following assumptions about the matching algorithm:

1. The algorithm itself is not specified; each SNO is free to use and tune any algorithm that meets the below criteria. Two of our prototype sites have made their matching algorithms available as part of this release, Indianapolis and Mendocino County. Boston uses a commercially available product, as do many existing health care systems.
2. Authorized queriers present a set of demographic details, and receive in return zero or more matching record locations. Only records meeting a minimum level of probability should be returned. That minimum level is calculated at each RLS such that the probability of returning a false positive match is very low (e.g., one chance in 100,000). Matches approaching but not reaching that level (sometimes called "fuzzy matches") should not be returned to avoid incidental disclosures. Further, the querier should not be told which data elements do not match since that could encourage fishing. It is legitimate

¹⁹ New York State Identification and Intelligence System, <http://www.nist.gov/dads/HTML/nysiis.html>.

to suggest that the querier provide additional data fields if these were not provided in the initial query. The details of how those matches are calculated must be hidden from the querier by the RLS, to preserve the ability of SNOs to use different and selectively tuned matching algorithms while maintaining standard interfaces.

3. Should the algorithm use transformations of the presented demographic data (e.g., treating Maggie and Margaret or off-by-one errors in numerical data as approximate matches) then the data returned should indicate both the fact of the match *and* the fact that a transformation was used in the match. The details of the display are up to the receiving application, but the information is provided to allow the requester, possibly in conversation with the patient, to add a check step against false positives, which are possible even with a high probability match.
4. Because delivering too little information is far less dangerous than delivering the wrong information, the algorithm must be tuned to minimize false matches, even at the expense of increasing the number of failed matches (false negatives). The algorithm must meet the policy requirements for accuracy, currently described in "Correctly Matching Patients with Their Records."
5. A national health identification number is not required. Demographic matching can work at population scale, without triggering either the enormous expense or political risk of failure that will attend any work on unique patient IDs. Should such an identifier exist, however, its use would still require the mechanics for matching and record location created by the RLS.²⁰ Social security number, although far from perfect as an identifier, and other types of identifying numbers, can increase the probability of achieving a correct match.
6. Individual SNOs may allow or require local IDs to be used as identifiers for the RLS (e.g., a SNO in a region with a primary employer may add employer IDs to the criteria to be matched.)
7. If there are records in the RLS that are below the matching threshold, the querier may not be presented a list to choose from, as this would create the very incidental disclosure the algorithm must be designed to avoid. (This restriction also forbids "wild card" searching, disallowing a search for e.g., all patients with the last name Smith.) Instead, the querier may be offered the ability to provide additional demographic details.
8. The RLS cannot assure that all records that exist for a given patient will be located, even in principal, because the patient may have received care outside the SNO, and because even within the SNO, there may be records that refer to the patient but are beneath the matching threshold, or that are being kept confidential for reasons of patient preference or legal constraints from the State or local policies set by the SNO or participating entities. The SNO may, at its discretion, require that displays of results returned from the RLS contain a reminder that the data may only be partial.

ISB

The other application a SNO is required to host is the Inter-SNO Bridge (ISB.) The ISB is the interface to data held by a SNO but used by institutions outside the SNO. It serves as a single point of access for all remote queries to entities inside any given SNO. The technical details of the ISB are in the "Health Information Exchange: Architecture Implementation Guide." The relevant policies are contained in *The **Connecting for Health** Common Framework: Resources for Implementing Private and Secure Health Information Exchange*.

The Common Framework makes the following assumptions about the design of the ISB:

²⁰ For a fuller accounting of **Connecting for Health's** view of national health identifiers, see the **Connecting for Health** 2005 report "Linking Health Care Information: Proposed Methods For Improving Care And Protecting Privacy," pp. 11-15. (http://connectingforhealth.org/assets/reports/linking_report_2_2005.pdf).

1. There is one ISB per SNO, which handles all per-patient clinical requests coming from or going to that SNO.
2. The ISB is *only* for patient-centered queries. Aggregate queries (e.g., "Find all admissions in the last 24 hours presenting with shortness of breath") should be dispatched to the participating institutions, or run against aggregated databases that are collected and kept offline. The lack of centralized clinical data keeps the ISB from being a target of loss or theft of clinical data, and allows interactions to be optimized for a single, simple case.
3. All transactions to and from the ISB are logged and audited.
4. The ISB must have a valid SSL certificate, and may only communicate with requestors who support encrypted web communications (https).
5. The ISB, like the RLS, is designed to take a query from authorized users in the form of demographic details or, alternatively, a query in the form of a unique provider ID plus a Medical Record Number.
6. The ISB *must* support patient data in incoming queries expressed in the HL7 2.4 format described in the "Health Information Exchange: Architecture Implementation Guide." The ISB *may* support incoming queries expressed in the HL7 3.0 format described in the "Health Information Exchange: Architecture Implementation Guide."
7. The ISB must support two possible patterns of request. The 'one pass' pattern has the requestor presenting patient details and receiving back the aggregated records. In this case, the ISB has acted as the aggregator, as described in Step 2 of the section *Modeling Interactions*, above.
8. The other pattern that must be supported is a 'two pass' interaction, in which the ISB acts like a standard RLS, returning locators to the remote querier, who then replies with a list of records they would like access to.
9. The ISB must support asynchronous delivery of records, where the requestor, whether a remote ISB or other entity, sends a request in, and then makes available a server for later delivery of the results of the request. "Later" may only be measured in seconds, but the asynchronous pattern is important because there is no guarantee that the ISB can dispatch and resolve all the required transactions local to the SNO quickly enough to support a synchronous return.

Part III: External Dependencies and Open Issues

Security: Identity, Authentication, and Authorization

Trust is the essential glue enabling transfer of medical records, and distrust is a key defense against their misuse. All the hardware and software in the world will not provide adequate defenses against users who are allowed to have copies of records from your institution, if those users fail to protect the records properly, or if they actively misuse those records. It is thus critical to provide the holder of any given set of records enough information to enable them to answer the question "Do I trust the person making this request, and the institution they work for?"

In a large network, it will be impossible for every possible pair of sender and receiver of a request to know one another in advance. Particularly for remote queries, some sort of transitive trust will be necessary -- "If I trust Institution A, I trust Institution A's employees." To make such a judgment, the requestor must be able to identify Institution A. In addition, the recipient of the request may need to review one or more queries, after the fact; it is essential, in this case, that the recipient be able to communicate quickly and effectively with the sender of the request. Being able to provide the time, date, entity identifier and person identifier will make such reviews considerably simpler, quicker and cheaper than merely having a transaction ID and requiring the counterparty to look it up in order to discover those identifiers on their own.

The federal HIPAA Privacy and Security Rules and applicable state laws provide the baseline for the **Connecting for Health** Common Framework, although in some cases greater privacy protections and individual rights are recommended by the CFH Policy Subcommittee. In no instance does the Common Framework permit less protection of personal health information than those required by federal or state law; however, participation in an SNO is not a surrogate for determining whether a Participant is a HIPAA Covered Entity or is in compliance with the HIPAA regulations. Importantly, the Common Framework permits SNO Participants to establish and follow their own more protective data management, privacy and security policies and procedures. In addition, some customization may be necessary at the SNO and Participant level to ensure consistency and compliance with applicable state laws. In addition to HIPAA, **Connecting for Health** recommends (but does not require) use of the use of the Common Criteria (ISO 15408) to analyze the security procedures of the SNO and of each participating entity.

In addition to HIPAA requirements, the transitive trust model requires reliable reporting of identity between parties participating in a record exchange, which in turn requires the issuing of identifiers for employees, and the maintenance of an authentication and authorization systems. Identifiers, authentication, and authorization are all required by HIPAA; what the Common Framework specifies are the reporting requirements necessary for transitive trust, plus the additional policies related to the transitive trust model, as laid out in "Authentication of System Users" in *The **Connecting for Health** Common Framework: Resources for Implementing Private and Secure Health Information Exchange*.

These reporting requirements have three levels:

1. Reporting of SNO identity for inter-SNO communications
2. Reporting of entity identity (the HIPAA-mandated identity)
3. Reporting of the identity of the individual who authorized the query.²¹

Reporting of SNO Identity: As noted above, the Common Framework assumes a 'thin NHIN'; the NHIN is a network of networks, with no NHIN-wide services required other than basic connectivity. Both the data and the contextual functions required for inter-SNO transfers of data rely on capabilities provided by the aggregate group of SNOs. There are no "top-level" service providers required.²² In particular, we do not envision a single top-level identity service. Instead, individual health care entities will use the identifiers mandated under HIPAA, and the SNOs are required to provide an invariant name, expressed as a URI,²³ that is unique within the NHIN.²⁴ In practice, the URI will also be the URL of the SNO, thus using the Domain Name System as the basic guarantor of uniqueness.

Reporting of Entity Identity: In the case of inter-SNO communications, it is not enough to know which SNO has dispatched a message; the source entity (e.g., clinic, hospital, lab) must also be identified. In practice, this identifier will be the HIPAA-

²¹ In the case of automatic queries, the individual identity should be that of the person most directly responsible for dispatching the query, e.g., the clerk who oversees the system doing the querying. This is because the goal of the identity reporting is to aid subsequent audits.

²² We do envision the possibility of third party providers offering network access to their services, but these services exist at the same conceptual layer as the SNOs themselves.

²³ Uniform Resource Identifier, which is effectively a location-insensitive version of a URL. For more on URIs and URLs, see <http://www.w3.org/TR/uri-clarification/>.

²⁴ Though guaranteed uniqueness presents interesting theoretical problems in large systems, it does not in small systems. Therefore, while we hope that someday the NHIN has enough participants to merit inclusion of one or more health care-specific schemes for mandated uniqueness, along the lines of the Internet's Domain Name System (DNS), we do not envision a system that large for 3 years at the earliest.

mandated National Provider Identifier (NPI),²⁵ plus a human-readable name of the institution.

Reporting of Individual Identity: In the case of inter-SNO communications, it is not enough to know which SNO and entity has dispatched a message; the requesting individual must also be identified. There is no person equivalent of the NPI under HIPAA; therefore some form of username or employee ID unique within the domain of a particular NPI must be reported, along with the human-readable name of the person making the request.

One possibility we explored but did not require was the additional reporting of the role of the requesting individual, e.g., clerk, admitting physician, etc. The reporting of such roles would be valuable in cases where the institution receiving the request wanted to limit responses to queriers with certain roles. However, the reporting of roles today is not universally adopted, and where it is in use, the definitions of the roles themselves are variable between institutions. If role-based access becomes important to the operation of the NHIN, considerable work will need to be done to standardize the expression of such roles.

As noted above, the reporting of these three identifiers -- SNO, entity, person -- serves two separate functions. The first is to allow judgment about whether to honor a given request for records, based on past expectations. The second is to simplify future audits should a particular transaction or set of transactions come under suspicion.

Based on the transmission of these identifiers, there are four conceptual levels of record exchange:

1. Within a single entity
2. Between entities in a SNO
3. Between entities in different SNOs
4. Requests from an entity not affiliated with a SNO

Within an Entity: Because this takes place beneath the level of SNO-wide traffic, these transactions are governed by the entity alone, and raise no SNO-wide identity reporting requirements.

Between Entities in a SNO: These transactions are governed by SNO-wide policies, but in all cases, the entity making a request, either to the RLS or to other entities, must report its NPI and the username and name of the person making the request. Because both sender and receiver are part of the same SNO, we presume that most of these requests will be routine, and will be subjected to a routine level of scrutiny.

Between Entities in Different SNOs: These transactions are not governed by any mutual contractual obligations, as entities in a single SNO are. Instead, they are governed by HIPAA and other national regulations and by the adoption of the Common Framework by each SNO. In all cases, the entity making the request must report the SNO of which it is a member, its NPI, and the username and name of the person making the request. Because sender and receiver of a query are not part of the same SNO, we assume these queries will be relatively rare (in comparison to the volume of local traffic), and should be subjected to a higher level of scrutiny.

²⁵ Because the NPI suffers from the same drawbacks as the SSN -- it is a public identifier with no accompanying authentication method -- the presentation of a HIPAA number must never be regarded as authenticating the requesting institution.

Requests from an entity not affiliated with a SNO:²⁶ The entity making a request to the SNO must report its NPI and the username and name of the person making the request. Because the sender is not part of the any SNO, we presume that most of all requests will be regarded as unusual, and will be subjected to the highest level of scrutiny.

Preventive Security: Encryption and Certificate Exchange

Preventive security encompasses defenses against outside attack, insider misuse of data, accidental loss or deletion, and so on. Outside the issues particular to identity, authentication and authorization, preventive security involves defending against access to data by unauthorized parties, misuse of data by authorized parties, unauthorized alteration of data, and accidental disclosure or deletion of data.

As noted above, the principal governing requirement for security is HIPAA, which describes administrative, physical, and technical safeguards required for securing data. In addition to HIPAA requirements and the policy requirements of Identity, Authentication, and Authorization, the Common Framework makes the following two specific security requirements:

- **SSL/TLS Encryption:** All traffic within and between SNOs will be encrypted using Secure Socket Layers 3.0 (SSL) or Transport Layer Security v 1.0 (TLS)
- **Exchange of SNO Certificates Between Pairs of SNOs:** Each ISB must have an SSL certificate, and any two ISBs planning to exchange data must each have the other's certificate. (This is implicit in the required asynchronous data exchange pattern between ISBs.)

Note that making each new SNO exchange credentials with all current SNOs will become problematic as the NHIN grows large.²⁷ We believe, however, that exchange of certificates between pairs of SNOs is the proper model for the NHIN in its early years, for two reasons: first, the number of SNOs (and therefore of ISBs needing certificates) will be small. Second, the trust of the participants in the system is a critical predictor of its success -- a system that is technically feasible but unpalatable to real-world institutions will simply fail.

Taken together, we believe that requiring that each new entrant announce itself to the existing members, and requiring some formal per-pair handshake, will not be technically onerous for a network containing even a hundred separate SNOs (something that would occur in 2008 at the earliest), and would be beneficial in terms of assuring participants that the network contained no unknown actors. It will also provide a better environment for actually watching trust develop, thus providing experience valuable to the design of any later-stage certificate brokering system.²⁸

Two proposed but not currently adopted questions around encryption merit further examination: the encryption of data as stored on a disk, and the special issues involved in handling SSNs.

²⁶ It is possible to imagine an entity in a SNO making an outbound request to an unaffiliated entity, but as the recipient would not have implemented the Common Framework, the question of standards and policies for such a transaction would be ad hoc.

²⁷ The Internet's Domain Name System (DNS), which we have studied as an example of a distributed system, grew out of the failure of such addressing schemes as the number of Internet nodes approached 1,000.

²⁸ Another reason for caution in this domain is that while generic federated identity systems have been under development for much of the last decade, no such systems have achieved widespread use, and most current multi-party identity systems are industry- or institution-specific. Because of the unique responsibilities of health care providers as stewards of data whose misuse can be both catastrophic and irremediable for the patient, we have generally erred on the side of accepting less efficiency in return for more safety.

With regards to encryption, there is considerable work being done on on-disk encryption, where the contents of every file are stored in an encrypted format, and are only able to be decrypted when running in a trusted environment (e.g., with proper user authentication.) This means that should hardware containing information on patients be physically stolen (as happened frequently in recent years), the material contained would not be available to the holders of the physical disk without their also having access to the password or any other required forms of identification. Though such systems are not so widely adopted today that we feel comfortable mandating their use as part of the Common Framework, the increasing threat of identity theft, coupled with the increasingly large aggregated data sets being generated by the health care industry, means that continued attention needs to be paid to the possibility of including such a requirement.

The use of SSNs also creates a special set of problems. A patient's SSN can be a valuable tool in improving match accuracy, and is in use in many medical record systems. However, the SSN is also a weak and poorly designed authentication token, and a key input to identity theft and other forms of personal intrusion. As a result there are good reasons to both use and not use SSN, and while there are health information networks in operation today which require its use, there are others that forbid it.

It is clear that the rise in identity theft is going to lead to some regulation of the use of SSNs, whether industry-by-industry, state-by-state, or on a national level. As a result, any SNO implementing an RLS must take special care in deciding whether to use SSN as a match field, and, if so, must take special care in protecting the SSN, never returning the SSN when a match is made, even if an SSN was submitted for a patient, so as to prevent fishing. Though use of the last 4 digits of the SSN is on the rise, this is simply re-creating the original difficulty. An attacker who has the last 4 digits of a full SSN, but only needs to report those 4 digits to gain access to material, has the same advantage as an attacker who has the full SSN when the full SSN is required.²⁹

Reactive Security: Auditing and Logging

In addition to preventive security, procedures need to be in place for detection of security breaches and other misuse of data. While detection is an important part of securing any system, it is especially vital in the health care system. Health is not money, and health care is not banking -- where many industries can tolerate a relatively high degree of inconvenience in their user transactions, health care cannot. Waiting a couple of days to get a new ATM card is an inconvenience; waiting a couple of hours to get data on a patient's allergies can be fatal.

Because of this, there is a presumption towards disclosure in cases where a physician reports a critical need for the data, a presumption reflected in, for example, Break the Glass modes of access. In addition to preventive measures against accidental or intentional misuse of data, maintaining a good log of the use of the system, and periodically auditing that log to look for negative events or patterns is critical.

While preventative security requires a high degree of standardization, so that everyone can, for example, use the same encryption scheme, reactive security requires less standardization. Instead, the ability to detect and react to events after the fact requires policies that set a minimum level of technical requirements, which requirements can be met in ways most compatible with the rest of a particular SNO's technological choices.

²⁹ We also examined the use of one-way hashes as ways of allowing the RLS to use SSN-like accuracy in matching records, without ever holding any patient's actual SSN. While the work on this method is impressive, the implementation overhead is large both in terms of original cost and in negative effect on speed of individual queries. Like mandated on-disk encryption, we believe this is an area that merits further study.

For the first phase of this work, we have specified two sets of policies regarding reactive security:

1. Logging and Auditing
2. Managing breaches

Logging and Auditing: It is essential that information about all relevant transactions that touch clinical data be logged, so that the system can be periodically audited, and so that, should there be misuse of data, the events leading up to that misuse can be examined, and it is essential that these logs be immutable, so that once written, they cannot be edited or deleted (in order to prevent sophisticated attackers from removing traces of their work.) The other essential constraint is that such logs not contain the full record being transmitted, so that the logs themselves do not become an alternate target for attackers looking for clinical information. The Common Framework policies regarding logging and auditing are contained in **Connecting for Health**, “Auditing Access to and Use of a Health Information Exchange.”

Managing Breaches: In addition to preventing breaches of identifying or clinical data where possible, every SNO needs a set of policies for reacting to such breaches where they occur, and these policies will necessarily differ between cases such as accidental disclosure, suffering from a successful external attack, or having an authorized employee misuse data. The Common Framework policies regarding breaches of information are contained in “Breaches of Confidential Health Information” in *The **Connecting for Health** Common Framework: Resources for Implementing Private and Secure Health Information Exchange*.

External Standards

Though the **Connecting for Health** prototype is obviously standards-based, rather than assuming homogeneity of underlying technology, we are principally a consumer of standards rather than a producer of them. Much of the prototype work was an attempt to *avoid* setting standards where possible, either by designing a system tolerant of multiple standards (as with the separation of clinical message contents from their delivery envelope) or by taking advantage of existing standards.

Many of the technical issues requiring standardization were solved by adopting Web Services standards; there are two cases where that has not been sufficient:³⁰

1. Coordinating asynchronous communication between SNOs
2. Making assertions about the individuals responsible for a request

In the case of coordinating asynchronous communications, the obvious standard to use is WS-Addressing.³¹ In the case of assertions about the identity of requesting individuals, the obvious standard to use is SAML 2.0.³² In both cases, the standard as proposed has not been widely adopted, nor has it been implemented in many key Web Services platforms. WS-Addressing is also not yet an official standard, despite two years of work.

³⁰ There may be only two such cases because we have intentionally restricted our work to a well-defined set of problems. As the work continues, we expect to encounter more of the issues that arise when a standard is required, and there are one or more unfinished or impractical candidates to choose from.

³¹ <http://www.w3.org/Submission/ws-addressing/>.

³² <http://www.oasis-open.org/committees/security/>.

Our solution in both cases has been to narrowly and provisionally adopt those parts of the two standards most compatible with our goals, but to make it clear in the documentation and in the headers of the messages that these are provisional standards and may change.

Beyond the issues having to do with Web Services generally, there are two additional sources of standards that any working system will require. The first, of course, is clinical standards. Though the model we have designed for requesting and receiving data is insensitive to what data is being carried -- lab results, medication lists, radiology reports, etc. -- the system as a whole requires good clinical standards, in order to be able to support increasing automation of record handling and exchange, decision support, and other clinical functions. Our goal here has been to design a system that is insensitive to what the actual standards are, so that when and as they arise, they can be used in this network model without significant re-engineering. It is our hope that the increased attention to health care IT will result in improved definition and adoption of clinical standards.

In the multi-region test of the proposed **Connecting for Health** architecture, we tested the exchange of both medication history and laboratory results. For the medication history standard, we adopted an XML serialization of the NCPDP standard, covered in the "Medication History Standards." The expression of the lab results were taken from the ELINCS standard, though our work revealed some areas where the proposed standards may need to be modified to cover a larger potential range of uses. The laboratory standards we used are documented in the "Laboratory Results Standards."

Medical history and laboratory results are essential parts of any working medical network, but they cover only a small fraction of the currently available data types, and as the health care industry continues to embrace IT, the number of data types will grow. As a result, **Connecting for Health** will continue to look to external designers and validators to guide our adoption of clinical standards.³³

The other required source of external standards is naming. Naming, which is to say the creation and assignment of identifiers, is in many ways more complex than standardization, since the two key requirements of naming -- uniqueness³⁴ and permanence³⁵ -- generally require sophisticated schemes for generating and maintaining those identifiers. The key concept in naming is the namespace; any given identifier needs to be unique in a particular namespace, which is the conceptual superset of all such identifiers. Any NHIN will require a namespace for all SNOs; each SNO will need to manage the namespace of its member entities; and each entity will need to manage the namespace of its employees and other individuals who may have access to the records IT system.

In all these cases, we have used existing namespaces where possible. For uniqueness of SNOs, we have relied on the fact that each SNO must register a URL, and that all URLs are globally unique; for entity naming, we have assumed the use of the HIPAA-mandated NPI; and for person identification within an entity, we have assumed that the HIPAA-required security

³³ Note that this does not presuppose the involvement of formal standards development organizations (SDOs) in all cases; there are many examples of proposed standards that failed to get any adoption in the field, as well as examples of de facto standards that were only blessed by standards bodies after the fact, if at all. Given **Connecting for Health's** focus on practical implementation and incremental development, we have tended to prefer de facto but unannointed standards to proposed but unadopted ones.

³⁴ Uniqueness in this case is contextual. A URL only needs to be disambiguated from another URL, and email address from another email address, and so on.

³⁵ Permanence is the characteristic of an identifier being unique in time. A permanent identifier, once issued, is never re-used to refer to anything else. A permanent identifier may stop being valid, when the thing it points to may disappear, but it will never point to anything else.

measures will guarantee the presence of unique identifiers for employees and others with system access.

We believe that these choices will take care of the basics of naming participant institutions and individuals for the early NHIN. However, there are a number of ways in which naming could become more complex over time. In particular, if our architecture succeeds, there will be a number of third parties who are not themselves covered entities under HIPAA, and will therefore lack an NPI, but will nevertheless need identification in the context of the NHIN (e.g., service providers for translation or off-site storage of clinical records.) Some way of assigning names compatible with the NPI may need to be designed. Our current solutions to naming should therefore be regarded as provisional, with future examination of additional requirements for naming being part of later phases of work.

Open Issues

The Common Framework is a work in progress; we have attempted to provide the starting technical and policy framework necessary for creating an interoperable nationwide health information network. Given the size of the task, however, such a framework necessarily leaves open technical issues, which will need to be addressed in subsequent iterations.

What follows is a brief list of such issues. Note that the list only addresses concerns that are largely or solely technological; issues such as financial structures or incentives, while vital, are out of the scope of this document.

Subscription to Data Sources

The basic transactions modeled here are request/response -- ask for and receive a record. There are obviously cases where an institution wants to subscribe to a particular data source, and receive updates when the data at that source is updated or modified.

As we have seen from the growth of http-based subscription models in other domains (e.g., the Atom syndication format³⁶), it is possible to model subscription models by using timed or triggered updates from the data host, or timed or triggered requests from the requestor. Working out which of these patterns are most appropriate in what situations is still to be done.

Aggregated Record Access

The basic transactions modeled here are requests for one individual's records. Use cases such as data quality and biosurveillance require the delivery of aggregated data from the source entities. Current programs have a mix of direct access (e.g., CDC's Biosense program) and hierarchical access (e.g., reports to State entities who in turn report to CDC.) While the interfaces required by the Common Framework provide a method for both direct and hierarchical aggregation of records, and that such aggregation does not require and should not use the RLS, it is not clear what role, if any, the SNO should play in providing or managing those interfaces, including especially whether the ISB should provide an interface for aggregate queries, or some other SNO-wide interface should be provided, or no such interface should be provided above the level of the individual entities.

This is an especially complex and important set of questions, as it involves balancing improved measurement and observation of the health care system with risks to patient privacy, and because many of the entities with the right to access aggregate data are not covered by

³⁶ <http://www.atomenabled.org/developers/syndication/>.

HIPPA, making the regulatory framework for assumptions about privacy and security more complex.

Federated Trust

Though we have adopted a transitive trust model (to trust Institution A means to trust their employees, and to trust that A manages identifiers, authentication, and authorization well), the growth of the NHIN will put that system under some strain. There is considerable work being done on federated trust models, where information and assertions about authentication and authorization are more fully described and communicated between parties. Future work on the NHIN will need to consider this work, and, if it proves viable, figure out how to implement it within the NHIN itself.

Patient Authentication and Access

One core goal of the Common Framework is to make it simpler for patients to gain access to their own records. We believe that the network as designed lowers the cost and raises the effectiveness of providing access to patient data; however, many hurdles still exist before there is widespread use of patient-facing access points to clinical data. A key hurdle is authentication. There is no obvious way to allow a patient to authenticate their own identity outside the context of a relationship with existing care providers; the hurdle this poses could potentially be overcome if there were a way to authenticate patients directly. Examination of possible direct patient authentication solutions may be a valuable area of inquiry. These issues are also addressed in "Patients Access to Their Own Health Information," part of *The **Connecting for Health** Common Framework: Resources for Implementing Private and Secure Health Information Exchange*.

Third Party Participation

There is currently no mechanism for allowing entities not covered under HIPPA direct access to the NHIN. Possible rationales for allowing such access would be third-party data cleaning or warehousing solutions; third party format translations, and secure repositories acting on behalf of the patient. Presently, any third party connected to the system must be a sub-contractor to a single entity or SNO, so that the legal liability remains with that entity. There may need to be a framework for allowing third party participation in a way that increases the flexibility of the network without decreasing security.

Multi-SNO Patient Lookup

There is no national interface for search in the system as envisioned. This is to preserve the partitioning of clinical data, so as to prevent major privacy spills. However, such a system also requires the requesting party to have enough information to know which remote SNO to query. As the number of participating SNOs grows large, this requirement could become problematic, but a single national interface for lookup could be a significant risk to privacy, by making fishing easier. As the network grows, understanding how to balance effectiveness with privacy protection in multi-SNO lookups will be critical.

National Services

The NHIN as currently envisioned has no national services separate from the ISB interfaces. With the spread of National Provider Identifier numbers, increased complexity of multi-SNO lookups, data translation requirements, etc, there may be some services best provisioned for the NHIN as a whole. Some mechanism for providing these services, probably by providing for

modified ISB-style interfaces, will need to be designed so that these services, should they be required, are broadly interoperable from launch.

Response Time and SLAs

The Common Framework Technical prototype was designed to demonstrate interoperability among diverse actors. Significant optimizations remain to be performed in order to reduce the systems response time, and profiling will be required before Service-Level Agreements can be implemented (agreements that promise, for example, sub-5 second response time for standard queries.)

Metadata and Request Filtering

The state of medical data is currently quite variable, and there is no obvious way to suddenly and dramatically increase the quality of that data, either by providing it in more structured formats or by increasing the accuracy of the fields. However, as the number of data sources grows, clinicians could find themselves moving from a world of little to no data to seeing a flood of data. Some way of pre-filtering the incoming data will be required, ideally by being able to specify relevant data types (e.g., medication history, radiology, etc), but such filtering will require relatively clean and well-formatted metadata, and must respect policy constraints such as the separation of demographic from clinical data in the system. Understanding how to improve metadata, how to integrate it in a way that respects policy constraints, and how to extract value from semi-structured data as it currently exists will be vital to future work on the NHIN.

Stemming Proliferation of Duplicate Records

As we move to a world where more data is shared, duplicate records will proliferate. Hospital A receives and stores patient data from Lab B; later, when Clinic C requests data on that patient from both A and B, C will get two copies of the lab data. Some method of identifying and rationalizing duplicate data will need to be provided as the number of data sources and frequency of requests increases.

Conclusion

We began the technical work for **Connecting for Health** in early 2004, and are very pleased to have arrived, in a little over two years, at a proof-of-concept system that demonstrates the core goal of the technical work to date:

It is possible to link disparate entities and regions together in a way that improves information sharing while protecting privacy, and to do so using mainly existing standards, across many hardware and software platforms, and at relatively low cost.

With three different regions involved in deploying the actual hardware and software required, we've seen that standards-based communications can provide secure links between different health care entities and networks, and that the deployment will not require massive hardware or software upgrades; will not require uniformity of hardware, software, or implementation strategies; and can be undertaken by different groups of programmers all working to an agreed set of standards.

As important as this test has been, however, we want to emphasize the modesty of the achievement relative to the enormity of the task. The actual deployment in the three regions is still early and there is insufficient usage to gauge its success. While lowering the cost of implementation and working to integrate existing IT investments where possible is necessary, it

is not sufficient. If we want nationwide adoption of interoperable health IT, there is considerable additional work to be done both on extending the technical work documented here and on problems outside the technical domain, such as improving incentives for interoperability, and raising data quality at the point of collection.

We are optimistic that by taking on the issues of packing and transporting data securely among diverse sets of participants, we can create an environment where it will be easier to understand the non-technical problems, and easier to imagine solutions for them. To reach this goal, however, we need feedback on our efforts to date. In 2006, **Connecting for Health** will continue to host public forums for discussing the issues relevant to improving health care, and we welcome any feedback you might have on this document, on the technical architecture this document covers, and on the related policies documented in the accompanying *The **Connecting for Health** Common Framework: Resources for Implementing Private and Secure Health Information Exchange*.

Acknowledgements

The working groups in the three regions of the **Connecting for Health** prototype and the Technical Subcommittee have worked for over two years to create a prototype of a decentralized, standards-based, and privacy-protecting architecture for the exchange of health records. During that time, we have been fortunate to work with respected experts in the fields of health and information technology, all of whom have contributed their time, energy, and expertise to the transition from a basic set of principles to a working prototype. Our consultants and volunteers have worked long hours in meetings and conference calls to negotiate high-level questions of architectural design and low-level details of particular technical protocols. We offer them our heartfelt thanks for taking on this journey with us, and look forward to the remaining work ahead.

In addition, we would like to offer special thanks to the working groups who took the conceptual technical model and instantiated it as running code: for the Massachusetts test, John Halamka, Greg DeBor, Gail Fournier, Vinod Muralidhar, and John Calladine; for the Indiana test, J. Marc Overhage, Clement McDonald, Lonnie Blevins, and Andrew Martin; and for the California test, Will Ross and Don Grodecki.

Finally, we must note that none of this work would have been possible without the leadership and inspiration of Clay Shirky, who encouraged us to turn theory into practice, and whose unmatched skills at navigating and then capturing each progressive phase of our work over the last two years allowed us to do so.

Connecting for Health Technology Subcommittee

Clay Shirky, New York University, (Chair)

Laura Adams, Rhode Island Quality Institute

Steve Adams, RMD Networks

William Braithwaite, MD, eHealth Initiative, (Co-Chair, Policy Subcommittee)

Deleys Brandman, First Consulting Group

Bryan Breen, Cisco Systems, Inc.

Sophia Chang, MD, MPH, California HealthCare Foundation

Art Davidson, MD, MSPH, Denver Public Health

Didi Davis, Eclipsys, Healthcare Information and Management Systems Society, and Integrating the Healthcare Enterprise

Greg DeBor, Computer Sciences Corporation

Lyman Dennis, Partnership HealthPlan of California, Healthcare Information and Management Systems Society, and Integrating the Healthcare Enterprise

George Eisenberg, IBM Corporation

David Epstein, IBM Corporation

Linda Fischetti*, RN, MS, Veterans Health Administration

Mark Frisse, MD, MBA, MSc, Vanderbilt Center for Better Health (Co-Chair, Policy Subcommittee)

Don Grodecki, Browsersoft, Inc.,

John Halamka, MD, CareGroup Healthcare System

Bob Hedgcock, Wisconsin Health Information Exchange

Noreen Hurley, Tufts Health Plan

Charles Jaffe, MD, PhD, Intel Corporation

Timothy Kenney, GE Healthcare

Josh Lemieux, Omnimedix Institute

J.P. Little, RxHub

Christopher Lindop, Eastman Kodak Company

David Lubinski, Microsoft Corporation

Janet Marchibroda, eHealth Initiative

Gregory Andre Marinkovich*, MD, FAAP LTC, Marine Corps, Office of Secretary of Defense/Health Affairs

Patrick McMahon, Microsoft Corporation

Omid Moghadam, Intel Corporation

Don Mon, PhD, American Health Information Management Association

Bruno Nardone, IBM Corporation

J. Marc Overhage, MD, PhD, Indiana Health Information Exchange; Indiana University School of Medicine, Regenstrief Institute for Healthcare

George Peredy, MD, Kaiser Permanente, HealthConnect

Nick Ragouzis, Enosis Group, LLC

Rick Ratliff, SureScripts

Jere Retzer, Oregon Health and Science University

Wes Rishel, Gartner Group

Barry Rhodes*, PhD, Center for Disease Control, United States Department of Health and Human Services

Scott Schumacher, PhD, Initiate Systems, Inc.

Raymond W. Scott, Axolotl Corporation

Don Simborg, MD, American Medical Informatics Association

Geoff Smith, Meditech

Jonathan Teich, MD, PhD, Healthvision

Micky Tripathi, Massachusetts eHealth Collaborative

Charlene Underwood, Healthcare Information and Management Systems Society, EHR Vendor Association

Karen Van Hentenryck, HL-7

Jukka Valkonen, California HealthCare Foundation

Cynthia Wark*, CAPT, United States Public Health Service Commissioned Corps, Centers for Medicare and Medicaid Services, United States Department of Health and Human Services

Jon White*, MD, Agency for Healthcare Research and Quality, United States Department of Health and Human Services

Scott Williams, MD, MPH, HealthInsight

Amy Zimmerman-Levitan, MPH, Rhode Island Department of Health

**Note: Federal employees participate in the Subcommittee but make no endorsement*